

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR LETTERS PATENT

Discontinuity Edge Overdraw

Inventor(s):

Hugues Hoppe
John M. Snyder
Pedro V. Sander
Steven J. Gortler

ATTORNEY'S DOCKET NO. MS1-732US

1 **TECHNICAL FIELD**

2 This invention relates to computer graphics, and more particularly, to
3 techniques for reducing aliasing along boundaries.

5 **BACKGROUND**

6 Aliasing is an important problem when rendering triangle or other
7 polygonal meshes. Computer graphics has long dealt with the issue of creating
8 discrete images without aliasing. For the hardware-accelerated triangle rendering
9 pipeline, four forms of aliasing can be identified:

- 10 (1) Aliasing within triangle interiors (undersampling of shading function).

11 One such example is aliasing due to texture undersampling, which can
12 be efficiently handled using mipmaps or higher-quality anisotropic
13 filtering. Other shading variations, like pinpoint specular highlights,
14 can exhibit high frequencies that are more difficult to predict and
15 bandlimit.

- 16 (2) Aliasing at triangle edges (appearance discontinuities).

17 (3) Aliasing among triangles (subpixel-sized triangles, also known as the
18 “small object” problem). This problem is partially helped by level-of-
19 detail control. Robust solution requires adequate supersampling or
20 analytic antialiasing.

- 21 (4) Aliasing at triangle intersections (where one triangle passes through
22 another). Static intersections can be preprocessed to yield explicit sharp
23 edges using polyhedral CSG. Dynamic intersections are difficult to
24 antialias without supersampling.

With current graphics hardware, a simple technique for reducing aliasing is to supersample and filter the output image. On current displays (desktop screens of approximately 1K×1K resolution), 2×2 supersampling reduces but does not eliminate aliasing. Of course, a finer supersampling resolution further reduces aliasing but rapidly becomes impractical.

One can implement 2×2 supersampling either by increasing the frame buffer resolution by a factor of four, or by accumulating four subpixel-offset images of the same scene for each frame. Both approaches are costly. The first requires four times the frame buffer memory and four times the fill-rate. The second requires four times the geometry processing, four times the fill-rate, and the addition of an accumulation buffer. The impact is that fill-rate-bound rendering becomes up to four times slower, and memory capacity is consumed that could otherwise be devoted to storing texture maps or caching geometry.

Much of the aliasing in current hardware rendering occurs along discontinuity edges. Perhaps most objectionable are the “crawling jaggies” that appear near discontinuity edges as a model moves. Fig. 1 shows a magnified triangle edge 100 to illustrate such “crawling jaggies”. The edge 100 is drawn on a grid 102 of pixels 104. Pixels touching the edge 100 are assigned to the triangle edge, resulting in a stair-stepped, jagged artifact 106. When the image moves, this jagged artifact appears to crawl along the edge, making them highly noticeable to an observer.

Such artifacts are perceptible even at high display resolutions where static spatial aliasing is less obvious, and are observable even with 2×2 supersampling. Since discontinuity edges typically cover only a small fraction of pixels, supersampling *every* pixel seems a brute-force solution.

1 Many general techniques to reduce aliasing have been used in computer
2 graphics, including uniform supersampling, adaptive supersampling, analytic
3 prefILTERING, and stochastic sampling. Adaptive supersampling attempts to focus
4 computation on troublesome areas such as discontinuity edges. However, adaptive
5 supersampling is difficult to make robust and implement in hardware. Prefiltering
6 approaches band-limit the continuous signal corresponding to a geometric
7 primitive (such as a constant-colored polygon fragment), before actually point-
8 sampling it. They require expensive visibility determinations over areas rather
9 than points. Stochastic sampling methods convert aliasing to less objectionable
10 noise (rather than “crawling jaggies”), but still require oversampling to acceptably
11 reduce aliasing artifacts.

12 Coverage bitmask approaches supersample only coverage rather than full r ,
13 g , b , z samples. These are effective at reducing artifacts at discontinuity edges, but
14 fail to eliminate aliasing at triangle intersections. Like traditional uniform
15 supersampling, they are brute-force solutions since a coverage bitmask must be
16 computed and stored at every pixel (typically 16-32 extra bits). Moreover, these
17 schemes maintain a list of fragments projecting onto each pixel.

18 The OpenGL® application program interface (API) introduced by SGI
19 offers a “polygon antialiasing” feature, available on some high-end graphics
20 workstations, that renders polygons with antialiased boundaries. It uses a special
21 blending mode (source_alpha_saturate) and only works when the polygons are
22 sorted front-to-back. A similar feature is also exposed in the DirectX® API from
23 Microsoft Corporation.

24 Another approach is to only antialias discontinuity edges. One technique is
25 to tag discontinuity edges and antialias them using prefILTERING convolution in a

1 scanline renderer. Another technique infers discontinuity edges in an aliased
2 image as a post-process. Pixels near discontinuities are then modified to account
3 for coverage of the inferred edges. This method gets confused at texture
4 discontinuities, ignores temporal aliasing, and is likely too expensive to perform at
5 interactive rates.

6 More recently, silhouette clipping has been proposed as a way to reduce
7 antialiasing along discontinuity edges. See, P. Sander, X. Gu, S. Gortler, H.
8 Hoppe, and J. Snyder, "Silhouette Clipping", SIGGRAPH 2000, pp. 327-334.
9 With silhouette clipping, a coarse mesh is clipped to the exact silhouette of a
10 detailed mesh using the stencil buffer. By transferring the stencil to the alpha
11 buffer and redrawing silhouette edges as antialiased lines, the *external* silhouette is
12 antialiased. However, this solution is limited to external silhouettes, and is not
13 more generally applied to reducing aliasing at both internal and external
14 silhouettes, as well as all discontinuity edges.

15 In another recent paper, the authors sketch a two-pass software rendering
16 approach for antialiasing silhouette edges. See, F. Sauer, O. Masclef, Y. Robert,
17 and P. Deltour, "Outcast: programming towards a design aesthetic", 1999 Game
18 Developers Conference, pp. 811-827. Unfortunately, the second pass bloats
19 foreground pixels near silhouettes by computing edge coverage at each pixel. The
20 method handles only silhouettes and detects these by exhaustive search. The
21 paper lacks details on how polygons are rasterized or how the two passes are
22 composited. U.S. Patent 6,005,580 to Donovan, which is entitled "Method and
23 Apparatus for Performing Post-Process Antialiasing of Polygon Edges", describes
24 a hardware approach that transfers the aliased frame buffer contents into texture
25 memory, and uses this texture to overdraw antialiased edges in a second pass.

1 Another solution proposed by M. Wimmer in his downloadable viewer
2 "View3DX software" is to overdraw antialiased lines. Unfortunately, this
3 approach failed without software sorting of all polygons.

4 The documentation for Microsoft's DirectX® software also mentions the
5 use of edge overdraw to achieve antialiasing. It explains that redrawing every
6 edge in a scene can work without introducing major artifacts, but it can also be
7 computationally expensive. In addition, it can be difficult to determine which
8 edges should be antialiased. The most important edges to redraw are those
9 between areas of very different color (for example, silhouette edges) or boundaries
10 between very different materials. Antialiasing the edge between two polygons of
11 roughly the same color will have no effect, yet is still computationally expensive.

12 Building upon this knowledge base, the inventors have developed an edge
13 overdraw approach that effectively reduces aliasing. They make the approach
14 practical by efficiently detecting and rendering just the discontinuity edges, and
15 introduce methods to maintain temporal smoothness, spatial consistency, and
16 spatial sharpness.

1 **SUMMARY**

2 Reduction of aliasing artifacts along discontinuity edges of a rendered
3 polygon mesh is achieved by overdrawing the edges as antialiased lines. The
4 discontinuity edges are oriented consistently and blended as they approach
5 silhouettes in the mesh to avoid popping at the edge, thereby achieving a temporal
6 smoothness at the silhouettes. This temporal smoothness is balanced with a
7 competing desire to maintain spatial sharpness by utilizing an asymmetric
8 blending technique. To further improve results, the discontinuity edges are sorted
9 by depth prior to overdrawing them. These processes are effective at reducing the
10 temporal artifacts known as “crawling jaggies”.

11 In the described implementation, a process is divided into two phases: a
12 preprocess phase and a runtime phase. During the preprocess phase, sharp edges
13 in the polygon mesh that are view independent are identified. Consistent
14 orientation is then assigned to the sharp edges. Additionally, a data structure used
15 to identify silhouette edges during the runtime phase is constructed.

16 During the runtime phase, the polygonal mesh is rendered and stored in a
17 frame buffer. Silhouette edges with respect to the given viewpoint are found using
18 the data structure constructed during the preprocess phase. The silhouette edges
19 and sharp edges together form the discontinuity edges. Shading is applied to the
20 discontinuity edges through an asymmetric blending technique that blends the
21 discontinuity edges while minimizing blurriness at the edges. The discontinuity
22 edges are then optionally sorted according to their z-value to produce a back-to-
23 front order. The edges are then rendered as antialiased lines in the back-to-front
24 order and written into the frame buffer, effectively superimposing the antialiased
25 lines on the original edges.

1

2 **BRIEF DESCRIPTION OF THE DRAWINGS**

3 Fig. 1 illustrates a conventional aliasing artifact, known as “crawling
4 jaggies”, that occurs at an edge of a rendered polygon mesh.

5 Fig. 2 is a block diagram of an exemplary graphics computing device that
6 implements a discontinuity edge overdraw process for reducing the crawling
7 jaggies.

8 Fig. 3 illustrates three graphical images to help demonstrate the effects of
9 the discontinuity edge overdraw process.

10 Fig. 4 is a flow diagram of the discontinuity edge overdraw process
11 implemented by the device of Fig. 2.

12 Fig. 5 illustrates a four-frame series of a moving graphical image to
13 demonstrate a “popping” effect.

14 Fig. 6 illustrates a four-frame series of the moving graphical image that is
15 generated with additional shading to reduce the “popping” effect.

16 Fig. 7 illustrates a graphical image rendered using symmetric blending
17 techniques at discontinuity edges to demonstrate a blurriness effect.

18 Fig. 8 illustrates the graphical image rendered using asymmetric blending
19 techniques at discontinuity edges to reduce the blurriness effect.

20 Fig. 9 illustrates three instances of the same frame of a graphical image,
21 where the first frame instance is produced without any orientation of the
22 discontinuity edges, the second frame instance is produced with per-path
23 orientation, and the third frame instance is produced using global heuristic
24 orientation.

1 **DETAILED DESCRIPTION**

2 Aliasing is an important problem when rendering polygon meshes, such as
3 triangle meshes. A major component of aliasing occurs along discontinuity edges
4 such as silhouettes, creases, and material boundaries. Reduction of aliasing
5 artifacts along discontinuity edges of a render polygon mesh is achieved by
6 overdrawing the edges as antialiased lines. The result of edge overdraw differs
7 from traditional antialiasing methods like supersampling in that one side of each
8 discontinuity edge is “bloated” by a fraction of a pixel.

9 As further refinements, the discontinuity edges are oriented consistently
10 and blended as they approach silhouettes in the mesh to avoid popping at the edge,
11 thereby achieving a temporal smoothness at the silhouettes. Unfortunately, edge
12 blending may produce blurriness. Thus, to balance temporal smoothness with
13 spatial sharpness, an asymmetric blending technique is employed. To further
14 improve results, the discontinuity edges are sorted by depth prior to overdrawing
15 them. These processes effectively reduce the “crawling jaggies” artifact.

16

17 **Graphics Computing Device**

18 Fig. 2 shows an exemplary graphics computing device 200 formed as a
19 personal computer implemented with a graphics card. The computing device 200
20 includes a central processing unit (CPU) 202, a system memory 204, a graphics
21 processing unit (GPU) 206, a frame buffer 208 that forms the video memory, and a
22 display 210 (e.g., flat panel display, CRT monitor, etc.). The system memory 204
23 is implemented as a variety of one or more types of volatile memory (e.g., RAM,
24 DRAM, etc.) and nonvolatile memory (e.g., ROM, Flash, EEPROM, disk, CD-
25 ROM, etc.). The graphics processing unit 206 and frame buffer 208 are typically

1 located on the graphics card, as illustrated by the dashed line separating card
2 components from computer components.

3 The CPU 202, system memory 204, graphics processing unit 206, frame
4 buffer 208, and display 210 are connected via a set of one or more buses 212, 214,
5 216, and 218. The buses may be integrated as one large system bus, or
6 implemented as separate buses. Separate buses are illustrated to pictorially
7 represent a graphics rendering pipeline from the CPU 202 to the display 210. The
8 buses may be any of a variety of bus architectures including a memory bus or
9 memory controller, a peripheral bus, an accelerated graphics port, and a processor
10 or local bus using any of a variety of bus architectures. By way of example, such
11 architectures can include an Industry Standard Architecture (ISA) bus, a Micro
12 Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a Video
13 Electronics Standards Association (VESA) local bus, and a Peripheral Component
14 Interconnects (PCI) bus also known as a Mezzanine bus.

15 A polygon mesh 220 is stored in system memory 204, or alternatively in
16 frame buffer 208 as illustrated by the dashed lines. This mesh 220 holds a
17 collection of polygons that are used to generate graphic images. In one
18 implementation, the polygon mesh 220 is composed of a collection of triangles.

19 One or more functional modules are illustrated within various components
20 to demonstrate the functions performed by the graphics computing device 200, as
21 well as the exemplary components that perform such functions. The CPU 202 has
22 a discontinuity edge (DE) detector 230 that is formed in hardware or software to
23 detect edges in a polygon mesh 220. The discontinuity edges are categorized into
24 two edge types: (1) "silhouette edges", which limit the extent of the projected

surface, and (2) "sharp edges", which mark shading discontinuities due to material boundaries or discontinuities in material attributes like normals and colors.

A silhouette tree structure 232 is shown stored in system memory 204. The silhouette tree 232 is a data structure that is used by the DE detector 230 to identify silhouette edges at each frame for a particular viewpoint. The silhouette search structure 232 stores all edges so that those that are silhouettes with respect to a given view may be found quickly. Once found, the extracted silhouette edges are collected and stored together with the sharp edges in a discontinuity edge (DE) list 234. The DE list can either be unsorted or sorted (as explained below in more detail).

The CPU 202 also includes an edge sorter 236 and a shade/blend module 238. The edge sorter 236 performs optional edge sorting tasks to sort the edges in the DE list 234 according to which edges are visible. The shade/blend module 238 shades and/or blends the edges. The DE detector 230, edge sorter 236, and shade/blend module 238 may be implemented in hardware within the CPU 202, or as software that is stored in system memory and executed on the CPU 202.

The GPU 206 includes a rendering module 240 and an overdrawing module 242, which can be implemented in hardware or a combination of hardware and software. The rendering module 240 renders the triangles from the mesh 232 and places the rendered images in the frame buffer 208.

Fig. 3 shows a series 300 of three images references as numbers 302, 304, and 306. The first image 302 illustrates the original rendered image with antialiasing along certain edges.

The overdrawing module 242 overdraws the discontinuity edges detected by the DE detector 230 as antialiased lines. The overdrawing module 242 exploits

existing hardware capable of rendering antialiased lines, which is well-known in the art of computer graphics. The overdrawn edges are illustrated as image 304 in Fig. 3. The z-buffer is used to resolve visibility between the mesh triangles and the overdrawn edges. The number of discontinuity edges is typically much smaller than the number of triangles or pixels, so the overall frame time overhead is small.

The overdrawn edges are effectively superimposed on the rendered image in the frame buffer 208 to significantly reduce the “crawling jaggies” caused by the aliasing along the edges of the rendered mesh. The final result is illustrated as image 306 in Fig. 3. Notice the reduction in jaggies from the original image 302, thereby improving rendering quality. In addition to reducing aliasing along the edges, the shade/blend module 238 helps maintain temporal smoothness, spatial consistency, and spatial sharpness. The rendered image is then scanned out to display 210.

Overdraw Method

Fig. 4 shows a process 400 for rendering polygonal meshes in a way that reduces aliasing along edges. The process 400 is implemented by the graphics computing device 200 and where appropriate, will be described with reference to various exemplary components that perform the various acts. The acts shown in the process can be implemented in hardware, software, or a combination of hardware and software.

The process 400 is divided into two phases: a preprocess phase and a runtime phase. The preprocess phase is performed at the CPU 202 in preparation to the

INSAI

2 runtime phase. The blocks in Fig. 4 are separated by a dashed line to differentiate
3 between preprocess operations and runtime operations.

4 Preprocess Phase

5 At block 402, the CPU 202 begins the preprocess phase by identifying
6 sharp edges (SE). Recall that discontinuity edges are a union of sharp edges
7 (which mark shading discontinuities) and silhouette edges (which limit the extent
8 of the projected surface). Because sharp edges demarcate shading discontinuities,
9 this set of edges is static and hence view independent. Therefore, sharp edges can
10 be collected during the preprocess phase, rather than during runtime, and are
11 subsequently overdrawn at every frame. The sharp edges are stored in the DE list
12 234. Fortunately, the number of sharp edges is typically a small fraction of the
13 total number of edges. Consistent orientations are then assigned to the sharp
14 edges.

15 Silhouette edges, in contrast, are based on the viewpoint. An edge is a
16 silhouette edge if one of its adjacent faces is frontfacing and the other backfacing.
17 For many meshes, the average number of silhouette edges per view is only $O(\sqrt{n})$,
18 where n is the number of mesh edges. So, typically only a small fraction of mesh
19 edges will be overdrawn as silhouette edges.

20 Collecting the silhouette edges can be done in a brute-force manner by
21 checking all mesh edges in $O(n)$ time. To accelerate this process, the CPU-based
22 DE detector 230 employs a fast silhouette extraction algorithm whose average
23 running time is proportional to the number of output silhouette edges. An
24 exemplary algorithm is described by Sander, et al., in "Silhouette Clipping",
25 SIGGRAPH 2000, pp. 327-334.

1 At block 404, the CPU 202 constructs the silhouette tree structure that is
2 used by the algorithm as a search hierarchy in which nodes represent clusters of
3 mesh edges. The silhouette tree 232 is a data structure stored in system memory
4 204. Then, for a given viewpoint at runtime, the DE detector 230 traverses the
5 hierarchy and is able to quickly skip entire subtrees that contain no silhouette
6 edges.

7

8 Runtime Phase

9 At block 406, the rendering module 240 of the graphics processing unit 206
10 renders an image as a triangle mesh (or other polygon mesh). For efficiency, it is
11 specified as a display list of triangle strips. The z-buffer is used to resolve
12 occlusion, and is saved for use during edge overdraw. The image is placed in the
13 frame buffer 208.

14 At block 408, the silhouette edge detector 230 of the CPU 202 identifies
15 the silhouettes for a particular view or scene using the silhouette search structure
16 232. Found silhouettes are stored in the DE list 234. For a closed object,
17 silhouette edges that are *concave* (having an outer dihedral angle ≤ 180 degrees)
18 are always occluded. Therefore, such concave edges need not be entered into the
19 search structure 232. This typically reduces the number of edges in the structure
20 by 40%. Furthermore, since sharp edges are typically overdrawn, they too are
21 omitted, resulting in an additional reduction of about 10%.

22 At block 410, the discontinuity edge detector 230 produces a set of
23 discontinuity edges as a union of the sharp edges identified in preprocessing and
24 the silhouettes identified during runtime, which are both stored in the DE list 234.

1 Occasionally, artifacts are produced when silhouettes lie in front of other
2 discontinuity paths. These artifacts can be removed by sorting the edges in back-
3 to-front order prior to rendering them during the subsequent overdraw operation.
4 At block 412, the edge sorter 236 may be used to sort the discontinuity edges in
5 the DE list 234 according to visibility.

6 In one implementation, the edge sorter uses the z-buffer values that
7 determine the depth dimension of the edges. The edges are sorted according to the
8 distance from the viewpoint to the edge midpoint. Although this midpoint depth-
9 sort heuristic occasionally gives an incorrect sort, artifacts are rare and comprise
10 only a few isolated pixels. By comparison, traditional back-to-front *polygon*
11 rendering requires correct occlusion-based ordering since mistakes there are much
12 more evident.

13 Of all sharp edges, only the sharp edges on or near the silhouette occlude
14 other discontinuity edges. An edge is declared to be near the silhouette if it has β
15 < 0.1 or $\beta > 0.9$, where β is a blend parameter defined below. Thus, one
16 optimization is to include only these sharp edges in the sort along with the other
17 silhouette edges. The remaining sharp edges are simply drawn first.

18 The sorting operation incurs some computational cost, and hence is
19 optional. It is beneficial when there are many discontinuity edge crossings.

20 At block 414, the shade/blend module 238 of the CPU 202 shades and/or
21 blends the discontinuity edges. To shade each discontinuity edge, the module 238
22 uses shading parameters (e.g. normals, colors, textures, texture coordinates) taken
23 from the edge's neighboring faces, denoted as the *left* and *right* faces. How the
24 shading parameters of the left and right face are combined depends on the
25 category of the discontinuity edge.

1 First, consider silhouette edges. The case of a non-sharp silhouette edge is
2 relatively straight forward since the shading parameters of the two adjacent faces
3 agree. At a sharp silhouette edge, the shading parameters of the two faces are
4 different, and the edge should be shaded using the parameters of the *frontfacing*
5 adjacent face. Note that depending on object orientation, a given sharp edge may
6 appear on the silhouette with either the left face frontfacing, or the right face
7 frontfacing.

8 The troublesome case is that of a sharp edge not on the silhouette. To
9 maintain temporal continuity, the edge should somehow smoothly transition to the
10 shading parameters of either the left face or the right face as it approaches the
11 silhouette. Otherwise, abruptly switching the shading parameters from one face to
12 the other would result in a “popping” artifact.

13 Fig. 5 shows a four-frame series consisting of four frames 500(1)-500(4)
14 that are drawn without the benefit of edge blending to demonstrate the “popping”
15 effect. Notice that as the sharp edge approaches the silhouette with the wrong
16 shading from frame 500(3) to frame 500(4), a popping artifact is produced as the
17 darker top line instantly disappears.

18 To solve this problem, for intermediate views where both adjacent faces are
19 frontfacing, the shade/blend module 238 shades the edge as a combination of the
20 two faces’ shading states. A blend parameter β is computed based on the inner
21 products of the viewing direction with the two adjacent face normals:

22

23 $V = \text{eye} - \text{edge.midpoint}$

24 $\text{dotL} = V \cdot \text{edge.leftFace.faceNormal}$

25 $\text{dotR} = V \cdot \text{edge.rightFace.faceNormal}$

1 $\beta = \text{dotR} / (\text{dotL} + \text{dotR}).$

2
3 Shading is then blended using the following relationship:

4
5 $(1 - \beta) \text{leftShading} + (\beta) \text{rightShading}.$

6
7 To achieve this blending, two alternate schemes may be used: "blended-
8 draw" and "double-draw". The edge blended-draw scheme renders the edge once,
9 as a blended combination of the two shading functions. Ideally, the blending is
10 performed with post-shaded color values. For texture-mapped meshes, this is
11 achieved using hardware multitexturing to blend the two adjacent textures. For
12 Gouraud-shaded surfaces, current hardware does not permit blending of post-
13 shaded results (without resorting to shading on the host CPU). Future hardware
14 supporting programmable shading will permit post-shaded blending. For now, the
15 shading attributes (e.g. normals and colors) are interpolated prior to hardware
16 shading. One drawback is that blending of normals can cause false highlights on
17 sharp crease edges.

18 The edge double-draw scheme renders the antialiased edge twice, once
19 using the shading function of the left face, and once using that of the right face.
20 An opacity value (alpha) is specified for compositing each edge "over" the frame
21 buffer. At least one of the edge renderings must use alpha=1 to prevent the aliased
22 background pixels from showing through. Moreover, the backface shading must
23 be attenuated to zero as the edge approaches the silhouette, to avoid popping. If
24 this backface shading edge is the one drawn with alpha=1, there is no way to
25 eliminate its contribution by rendering the second antialiased line over it (due to

the antialiased line's partial coverage). Accordingly, a simple order-switching algorithm is used. Specifically, if $\beta < .5$, we first render with left face shading and alpha=1, followed by right face shading and alpha= β . Otherwise we first render with right face shading and alpha=1 followed by left face shading with alpha=1- β . Although this results in a slight discontinuity at the $\beta = 0.5$ transition, it is not perceptible in practice.

For blending, the edge double-draw scheme is preferred for its implementation simplicity. Also, it does not require multitexturing and does not exhibit false highlights due to pre-shaded blending.

Fig. 6 shows the same four-frame series as that of Fig. 5, but with the image in each frame being drawn using the double-draw scheme. Notice the smoother transition from frame-to-frame in comparison to non-blended images of Fig. 5.

Although blending is beneficial to avoid temporal popping, it tends to blur the discontinuity edge because the shading of the blended edge agrees with neither of the adjacent faces. To compromise between the competing goals of temporal smoothness and spatial sharpness, a hybrid approach uses the parameters from a single face (the *left* face) as much as possible, while still avoiding objectionable pops.

To improve the sharpness at the edge, the blend parameter β is mapped through an asymmetric transfer function:

$$\beta' = \begin{cases} 0 & , \text{ if } \beta \leq \tau \\ (\beta - \tau)/(1 - \tau) & , \text{ otherwise,} \end{cases}$$

1 The edges are blended using the resulting β' . We find that with τ set to 0.9,
2 edge transition are still temporally smooth, but the fraction of blended sharp edges
3 drops from about 30% to 2% on average. In addition to restoring edge sharpness
4 and saving blending operations, asymmetric blending allows most of the edge
5 geometry to remain static, possibly cached on the graphics card. This may be
6 exploited by first rendering all sharp edges as a display list and then the few
7 blended edges.

8 Figs. 7 and 8 show the same image, but Fig. 7 illustrates the image with
9 symmetric blending whereas Fig. 8 shows the image with asymmetric blending.
10 With asymmetric blending (Fig. 8), a non-silhouette sharp edge is usually drawn
11 using the shading parameters of the left face. This has the drawback of shifting
12 the proper material boundary by half a pixel. However, this is less objectionable
13 than the extra softening or blurring that occurs when using symmetric blending.

14 When referring earlier to the left/right faces of an edge, an edge orientation
15 is assumed. If a shading discontinuity consists of several sharp edges along a path
16 and the orientation of each edge in the path is selected independently, then an
17 asymmetric blending bias results in staggered-looking discontinuities. The
18 solution to reducing these discontinuities is to orient discontinuity edges
19 consistently using an orientation algorithm. Orientation of sharp edges may be
20 performed during the preprocessing phase prior to runtime.

21 Fig. 9 shows three instances 900(1)-900(3) of the same frame to illustrate
22 different orientation approaches. The first instance 900(1) illustrates the case
23 without orientation of the discontinuity edges to demonstrate the staggered-
24 looking discontinuities.

Using a per-path orientation approach, sharp edges are first concatenated together into sharp paths. Two adjacent sharp edges are placed in the same path if their shared vertex has no other adjacent sharp edges. For each path, we assign an orientation to one edge, and then locally propagate this orientation along the entire path. However, if each sharp path is oriented independently, some regular structures appear non-uniform. For example, the second frame instance 900(2) in Fig. 9 shows that some patches might appear larger than others.

This is resolved using a global heuristic that begins by picking two arbitrary orthogonal vectors, such as $g_1=(2,5,1)$ and $g_2=(2,1,-9)$. For each sharp path, a representative edge is determined as the one whose midpoint is farthest along the vector g_1 . The orientation of this edge is then assigned based on the sign of the dot product between the edge vector and the vector g_2 . Given this first edge orientation, the process locally propagates along the sharp path as before. The result is shown as the third frame instance 900(3) in Fig. 9.

With reference again to Fig. 4, at block 416, the overdrawing module 242 of the graphics processing unit 206 overdraws the discontinuity edges with antialiasing lines. In one implementation, once shading parameters are determined, the edges are once again rendered into the frame buffer 208 as antialiased lines. Alpha blending is configured so that the lines are drawn using an “over” operation. The z-buffer test is enabled to avoid drawing occluded edges. The z-buffer write is disabled so that chains of antialiased edges do not have gaps at the shared endpoints between individual edges.

The rendering process 400 may be implemented in software that, when executed, directs the graphics computing device to perform the various functions.

1 The following pseudo code represents an exemplary implementation of the process
2 400.

3 **Preprocess**

4 Collect sharp edges *Sharp* in scene;
5 Assign consistent orientations to *Sharp*;
6 Construct silhouette extraction tree (excluding sharp & concave);

7 **Runtime (given viewpoint for each frame)**

8 Render scene;

9 $S = \emptyset$;

10 **for** edge *e* in *Sharp*

11 $d_{left} = \text{dot}(e.f_{left}.normal, e.midpoint - viewpoint);$
12 $d_{right} = \text{dot}(e.f_{right}.normal, e.midpoint - viewpoint);$
13 **if** $d_{left} < 0$ **and** $d_{right} < 0$ **then continue**; // backfacing
14 $e.\beta = d_{right} / (d_{left} + d_{right});$
15 **if** $0.1 < e.\beta < 0.9$ **then**

16 Render *e* with $\alpha = 1.0$ using *e.f_{left}* shading;

17 **else**

18 $S = S \cup \{e\};$

19 Extract silhouette edges *Sil* given viewpoint,

20 $S = S \cup Sil;$

21 Sort *S* in back-to-front order;

22 **for** edge *e* in *S*

23 **if** *e* $\in Sil$ **then**

24 Render *e* with $\alpha = 1.0$ using *e.f_{front}* shading;

25 **else if** $e.\beta < 0.9$ **then**

26 Render *e* with $\alpha = 1.0$ using *e.f_{left}* shading;

27 **else**

28 $e.\beta' = (e.\beta - 0.9) / (1.0 - 0.9);$

29 **if** $e.\beta' < 0.5$ **then**

30 Render *e* with $\alpha = 1.0$ using *e.f_{left}* shading;

31 Render *e* with $\alpha = e.\beta'$ using *e.f_{right}* shading;

32 **else**

33 Render *e* with $\alpha = 1.0$ using *e.f_{right}* shading;

34 Render *e* with $\alpha = 1.0 - e.\beta'$ using *e.f_{left}* shading;

1 **Experimental Results**

2 The graphics rendering process described above was implemented, tested,
3 and observed by the inventors. The software was written using OpenGL. It was
4 implemented and tested on a Pentium III 800MHz PC with an NVIDIA GeForce2
5 graphics card. The inventors further verified that the process works on an SGI
6 Octane graphics card.

7 The inventors used the OpenGL® API function `glEnable(GL_POLYGON_`
8 `OFFSET_FILL)` to perturb z-buffer values of triangles behind those of lines. This
9 was done so that antialiased lines pass the z-buffer test to cover the crawling
10 jaggies. For edges adjacent to triangles with high depth slope, the inventors
11 observed remaining aliasing artifacts, suggesting that the function
12 `glPolygonOffset()` feature is not pushing the triangles back sufficiently. The
13 presence of these artifacts varies with the particular graphics hardware.

14 For efficiency, `GL_BLEND` is enabled for rendering lines. The lines are
15 rendered using the default `glLineWidth(1.0f)`.

16 When edge sorting is enabled, `qsort()` is employed. A faster algorithm,
17 like bucket sort, could further improve the timing results when rendering high-
18 resolution models.

19 The system was tested on six models. The preprocessing bottleneck is the
20 creation of the silhouette tree, which is currently unoptimized and can take several
21 minutes on large models. Collecting the sharp edges and assigning them
22 consistent orientations takes only a few seconds.

23 Runtime results are shown in Table 1. Note that the extracted silhouette
24 edges do not include silhouette edges that are sharp or concave. Rendered edges
25 excludes backfacing sharp edges. The ship example has a higher performance

overhead because it is geometry-bound and has a high number of discontinuity edges.

Model	man	plane	stoneh	dino	ship
<i>Faces</i>	1,586	8,000	1,380	43,866	85,068
<i>Edges</i>	2,379	12,000	2,070	65,799	127,602
<i>Sharp edges</i>	354	2,085	1,250	900	19,769
<i>Edge statistics averaged over 100 viewpoints</i>					
<i>Extracted sil. edges</i>	94	393	22	365	7,122
<i>Rendered edges</i>	373	1,727	952	1,894	21,980
<i>Sorted edges</i>	309	1,212	661	1,240	16,448
<i>Blended edges</i>	6	23	10	23	266
<i>Rendering time per frame (in milliseconds)</i>					
<i>No edge overdraw</i>	7.2	9.8	9.6	18.9	40.1
<i>Unsorted edge overdraw</i>	7.7	10.3	10.7	20.0	88.4
<i>Sorted edge overdraw</i>	7.7	10.8	10.7	23.3	121.2

Table 1

Conclusion

Described above is an edge overdraw technique that reduces discontinuity edge artifacts. For typical models having a small proportion of discontinuity edges, edge overdraw can be performed with little added cost. While the method is designed for spatial antialiasing, it reduces the “crawling jaggies” artifact.

The process generalizes to the case of meshes with boundaries. A boundary edge can be thought of as a smooth edge with an outer dihedral angle of

1 360 degrees. Thus, it is reported as a silhouette edge for all viewpoints. The edge
2 is shaded using the attributes of its one adjacent face. With surface boundaries,
3 the mesh interior may become visible, so some of the optimizations may be
4 disabled. For instance, concave edges are no longer omitted from the silhouette
5 search structure and sharp edges are drawn even if they are backfacing.

6 Edge overdraw extends triangles by a fraction of a pixel along
7 discontinuities. At silhouette edges, this essentially enlarges the foreground object
8 slightly at the expense of the background. This is done since the frame buffer
9 lacks information about what lies behind the foreground object at partially covered
10 pixels drawn in the foreground.

11 For a scene with many objects, edges can be overdrawn after all objects are
12 rendered. Alternatively, edge overdraw can be applied after the rendering of each
13 object. In that case, the objects are rendered in back-to-front order if one desires
14 correct behavior at object silhouettes.

15 Although the invention has been described in language specific to structural
16 features and/or methodological acts, it is to be understood that the invention
17 defined in the appended claims is not necessarily limited to the specific features or
18 acts described. Rather, the specific features and acts are disclosed as exemplary
19 forms of implementing the claimed invention.

20

21

22

23

24

25